

Max-Throughput for (Conservative) k -of- n Testing

Lisa Hellerstein*

Özgür Özkan†

Linda Sellie‡

September 16, 2011

Abstract

We define a variant of k -of- n testing that we call *conservative* k -of- n testing. We present a polynomial-time, combinatorial algorithm for the problem of maximizing throughput of conservative k -of- n testing, in a parallel setting. This extends previous work of Kodialam and Condon et al., who presented combinatorial algorithms for parallel pipelined filter ordering, which is the special case where $k = 1$ (or $k = n$) [7, 3, 4]. We also consider the problem of maximizing throughput for *standard* k -of- n testing, and show how to obtain a polynomial-time algorithm based on the ellipsoid method using previous techniques.

1 Introduction

In *standard* k -of- n testing, there are n binary tests, that can be applied to an “item” x . We use x_i to denote the value of the i^{th} test on x , and treat x as an element of $\{0, 1\}^n$. With probability p_i , $x_i = 1$, and with probability $1 - p_i$, $x_i = 0$. The tests are independent, and we are given p_1, \dots, p_n . We need to determine whether at least k of the n tests on x have a value of 0, by applying the tests sequentially to x . Once we have enough information to determine whether this is the case, that is, *once we have observed k tests with value 0, or $n - k + 1$ tests with value 1*, we do not need to perform further tests.¹

We define *conservative* k -of- n testing the same way, except that we continue performing tests until we have either observed k tests with value 0, or have performed all n tests. In particular, we do not stop testing when we have observed $n - k + 1$ tests with value 1.

There are many applications where k -of- n testing problems arise, including quality testing, medical diagnosis, and database query optimization. In quality testing, an item x manufactured by a factory is tested for defects. If it has at least k defects, it is discarded. In medical diagnosis, the item x is a patient; patients are diagnosed with a particular disease if they fail at least k out of n special medical tests. A database query may ask for all tuples x satisfying at least k of n given predicates (typically $k = 1$ or $k = n$).

For $k = 1$, standard and conservative k -of- n testing are the same. For $k > 1$, the conservative variant is relevant in a setting where, for items failing fewer than k tests, we need to know *which* tests they failed. For example, in quality testing, we may want to know which tests were failed by items failing fewer than k tests (i.e. those not discarded) in order to repair the associated defects.

Our focus is on the MAXTHROUGHPUT problem for k -of- n testing. Here the objective is to maximize the throughput of a system for k -of- n testing in a parallel setting where each test is performed by a separate “processor”. In this problem, in addition to the probabilities p_i , there is a *rate limit* r_i associated with the processor that performs test i , indicating that the processor can only perform tests on r_i items per unit time.

MAXTHROUGHPUT problems are closely related to MINCOST problems [8, 5]. In the MINCOST problem for k -of- n testing, in addition to the probabilities p_i , there is a cost c_i associated with performing the i^{th}

*This research is supported by the NSF Grant CCF-0917153.

†This research supported by US Department of Education grant P200A090157.

‡This research is supported by a CIFellows Project postdoc, sponsored by NSF and the CRA.

¹In an alternative definition of k -of- n testing, the task is to determine whether at least k of the n tests have a value of 1. Symmetric results hold for this definition.

test. The goal is to find a testing strategy (i.e. decision tree) that minimizes the expected cost of testing an individual item. There are polynomial-time algorithms for solving the MINCOST problem for standard k -of- n testing [9, 10, 1, 2].

Kodialam was the first to study the MAXTHROUGHPUT k -of- n testing problem, in the special case where $k = 1$ [7]. He gave a $\mathcal{O}(n^3 \log n)$ algorithm for the problem. The algorithm is combinatorial, but its correctness proof relies on polymatroid theory. Later, Condon et. al. studied the problem, calling it “parallel pipelined filter ordering”. They gave a $\mathcal{O}(n^2)$ combinatorial algorithm, with a direct correctness proof [4].

In this paper, we extend the previous work by giving a polynomial-time combinatorial algorithm for the MAXTHROUGHPUT problem for conservative k -of- n testing. Our algorithm can be implemented using simple dynamic programming to run in time $\mathcal{O}(kn^2)$, which is $\mathcal{O}(n^2)$ for constant k . We leave for future work the problem of reducing the runtime for non-constant k .

The MAXTHROUGHPUT problem for standard k -of- n testing appears to be fundamentally different from its conservative variant. We leave as an open problem the task of developing a polynomial time *combinatorial* algorithm for this problem. We show that previous techniques can be used to obtain a polynomial-time algorithm based on the ellipsoid method. This approach could also be used to yield an algorithm, based on the ellipsoid method, for the conservative variant.

2 Related Work

Deshpande and Hellerstein studied the MAXTHROUGHPUT problem for $k = 1$, when there are precedence constraints between tests [5]. They also showed a close relationship between the exact MINCOST and MAXTHROUGHPUT problems for k -of- n testing, when $k = 1$. Their results can be generalized to apply to testing of other functions.

Liu et al. [8] presented a generic, linear-programming based, method for converting an approximation algorithm for a MINCOST problem, into an approximation algorithm for a MAXTHROUGHPUT problem. Their results are not applicable to this paper, where we consider only exact algorithms.

Polynomial-time algorithms for the MINCOST problem for standard k -of- n testing were given by Salloum, Breuer, Ben-Dov, and Chang et al. [9, 10, 1, 2, 11].

The problem of how to best order a sequence of tests, in a sequential setting, has been studied in many different contexts, and in many different models. See e.g. [8] and [4] for a discussion of related work on the filter-ordering problem (i.e. the MINCOST problem for $k = 1$) and its variants, and [12] for a general survey of sequential testing of functions.

3 Problem Definitions

A k -of- n testing strategy for tests $1, \dots, n$ is a binary decision tree T that computes the k -of- n function, $f : \{0, 1\}^n \rightarrow \{0, 1\}$, where $f(x) = 1$ if and only if x contains fewer than k 0’s. Each node of T is labeled by a variable x_i . The left child of a node labeled with x_i is associated with $x_i = 0$ (i.e., failing test i), and the right child with $x_i = 1$ (i.e., passing test i). Each $x \in \{0, 1\}^n$ corresponds to a root-to-leaf path in the usual way, and the label at the leaf is $f(x)$.

A k -of- n testing strategy T is *conservative* if, for each root-to-leaf path leading to a leaf labeled 1, the path contains exactly n non-leaf nodes, each labeled with a distinct variable x_i .

Given a permutation π of the n tests, we define $T_k^c(\pi)$ to be the conservative strategy described by the following procedure: *Perform the tests in order of permutation π until at least k 0’s have been observed, or all tests have been performed, whichever comes first. Output 0 in the first case, and 1 in the second.*

Similarly, we define $T_k^s(\pi)$ to be the following standard k -of- n testing strategy: *Perform the tests in order of permutation π until at least k 0’s have been observed, or until $n - k + 1$ 1’s have been observed, whichever comes first. Output 0 in the first case, and 1 in the second.*

Each test i has an associated probability p_i , where $0 < p_i < 1$. Let D_p denote the product distribution on $\{0, 1\}^n$ defined by the p_i ’s; that is, if x is drawn from D_p , then $\forall i, \mathbf{Pr}[x_i = 1] = p_i$ and the x_i are independent.

We use $x \sim D_p$ to denote a random x drawn from D_p . In what follows, when we use an expression of the form $\text{Prob}[\dots]$ involving an item x , we mean the probability with respect to D_p .

3.1 The MinCost problem

In the MINCOST problem for standard k -of- n testing, we are given n probabilities p_i and costs $c_i > 0$, for $i \in \{1, \dots, n\}$, associated with the tests. The goal is to find a k -of- n testing strategy T that minimizes the expected cost of applying T to a random item $x \sim D_p$. The cost of applying a testing strategy T to an item x is the sum of the costs of the tests along the root-to-leaf path for x in T .

In the MINCOST problem for conservative k -of- n testing, the goal is the same, except that we are restricted to finding a *conservative* testing strategy.

For example, consider the MINCOST 2-of-3 problem with probabilities $p_1 = p_2 = 1/2$, $p_3 = 1/3$ and costs $c_1 = 1$, $c_2 = c_3 = 2$. A standard testing strategy for this problem can be described procedurally as follows: *Given item x , begin by performing test 1. If $x_1 = 1$, follow strategy $T_2^s(\pi_1)$, where $\pi_1 = (2, 3)$. Else if $x_1 = 0$, follow strategy $T_1^s(\pi_2)$, where $\pi_2 = (3, 2)$.*

Under the above strategy, which can be shown to be optimal, evaluating $x = (0, 0, 1)$ costs 5, and evaluating $x' = (1, 1, 0)$ costs 3. The expected cost of applying this strategy to a random item $x \sim D_p$ is $3\frac{5}{6}$.

Because the MINCOST testing strategy may be a tree of exponential size in the number of tests, algorithms for the MINCOST problem may output a compact representation of the output strategy.

3.2 The MaxThroughput problem

The MAXTHROUGHPUT problem for k -of- n testing is a natural generalization of the MAXTHROUGHPUT problem for 1-of- n testing, first studied by Kodialam [7]. We give basic definitions and motivation here. For further information, see [7, 3, 4].

In the MAXTHROUGHPUT problem for k -of- n testing, as in the MINCOST problem, we are given the probabilities p_1, \dots, p_n associated with the tests. Instead of costs c_i for the tests, we are given *rate limits* $r_i > 0$. The MAXTHROUGHPUT problem arises in the following context. There is an (effectively infinite) stream of items x that need to be tested. Every item x must be assigned a strategy T that will determine which tests are performed on it. Different items may be assigned to different strategies. Each test is performed by a separate “processor”, and the processors operate in parallel. (Imagine a factory testing setting.) Item x is sent from processor to processor for testing, according to its strategy T . Each processor can only test one item at a time. We view the problem of assigning items to strategies as a flow-routing problem.

Processor O_i performs test i . It has rate limit (capacity) r_i , indicating that it can only process r_i items x per unit time.

The goal is to determine how many items should be assigned to each strategy T , per unit time, in order to maximize the number of items that can be processed per unit time, the throughput of the system. The solution must respect the rate limits of the processors, in that the expected number of items that need to be tested by processor O_i per unit time must not exceed r_i . We assume that tests behave according to expectation: if m items are tested by processor O_i per unit time, then mp_i of them will have the value 1, and $m(1 - p_i)$ will have the value 0.

Let \mathcal{T} denote the set of all k -of- n testing strategies and \mathcal{T}_c denote the set of all conservative k -of- n testing strategies. Formally, the MAXTHROUGHPUT problem for standard k -of- n testing is defined by the linear program below. The linear program defining the MAXTHROUGHPUT problem for conservative k -of- n testing is obtained by simply replacing the set of k -of- n testing strategies \mathcal{T} by the set of conservative k -of- n testing strategies \mathcal{T}_c .

We refer to a feasible assignment to the variables z_T in the below LP as a *routing*. We call constraints of type (1) *rate constraints*. The value of F is the *throughput* of the routing. For $i \in \{1, \dots, n\}$, if $\sum_{T \in \mathcal{T}} g(T, i)z_T = r_i$, we say that the routing *saturates* processor O_i .

We will refer to the MAXTHROUGHPUT problems for standard and conservative k -of- n testing as “SMT(k) problem” and “CMT(k) problem”, respectively.

As a simple example, consider the following $CMT(k)$ problem (equivalently, $SMT(k)$ problem) instance, where $k = 1$ and $n = 2$: $r_1 = 1$, $r_2 = 2$, $p_1 = 1/2$, $p_2 = 1/4$. There are only two possible strategies, $T_1(\pi_1)$, where $\pi_1 = (1, 2)$, and $T_1(\pi_2)$, where $\pi_2 = (2, 1)$. Since all flow assigned to $T_1(\pi_1)$ is tested by O_1 , $g(T_1(\pi_1), 1) = 1$; this flow continues on to O_2 only if it passes test 1, which happens with probability $p_1 = 1/2$, so $g(T_1(\pi_1), 2) = 1/2$. Similarly, $g(T_1(\pi_2), 2) = 1$ while $g(T_1(\pi_2), 1) = 1/4$, since $p_2 = 1/4$. Consider the routing that assigns $F_1 = 4/7$ units of flow to strategy $T_1(\pi_1)$, and $F_2 = 12/7$ units to strategy $T_1(\pi_2)$. Then the amount of flow reaching O_1 is $4/7 \times g(T_1(\pi_1), 1) + 12/7 \times g(T_1(\pi_2), 1) = 1$, and the amount of flow reaching O_2 is $4/7 \times g(T_1(\pi_1), 2) + 12/7 \times g(T_1(\pi_2), 2) = 2$. Since $r_1 = 1$ and $r_2 = 2$, this routing saturates both processors. By the results of Condon et al. [4], it is optimal.

MaxThroughput LP:

Given $r_1, \dots, r_n > 0$ and $p_1, \dots, p_n \in (0, 1)$, find an assignment to the variables z_T , for all $T \in \mathcal{T}$, that maximizes

$$F = \sum_{T \in \mathcal{T}} z_T$$

subject to the constraints:

- (1) $\sum_{T \in \mathcal{T}} g(T, i) z_T \leq r_i$ for all $i \in \{1, \dots, n\}$ and
- (2) $z_T \geq 0$ for all $T \in \mathcal{T}$

where $g(T, i)$ denotes the probability that test i will be performed on an item x that is tested using strategy T , when $x \sim D_p$.

4 The Algorithm for the MinCost Problem

In the literature, versions of the MINCOST problem for 1-of- n testing are studied under a variety of different names, including pipelined filter ordering, selection ordering, and satisficing search (cf. [4]).

The following is a well-known, simple algorithm for solving the MINCOST problem for standard 1-of- n testing (see e.g. [6]): First, sort the tests in increasing order of the ratio $c_i/(1 - p_i)$. Next, renumber the tests, so that $c_1/(1 - p_1) < c_2/(1 - p_2) < \dots < c_n/(1 - p_n)$. Finally, output the sorted list $\pi = (1, \dots, n)$ of tests, which is a compact representation of the strategy $T_1^s(\pi)$ (which is the same as $T_1^c(\pi)$).

The above algorithm can be applied to the MINCOST problem for conservative k -of- n testing, simply by treating π as a compact representation of the conservative strategy $T_k^c(\pi)$. We now show that $T_k^c(\pi)$ is, in fact, a MINCOST conservative strategy. The proof is a generalization of the proof of correctness for 1-of- n testing, and is slightly more complicated here in part because we need to consider the possibility that the optimal strategy corresponds to a complicated decision tree, rather than one specified by a single permutation.

Lemma 1. *The strategy $T_k^c(\pi)$ output by the above algorithm has minimum expected cost among all conservative strategies.*

Proof. Our proof is by induction on n . The base case, where $n = 1$ is trivially true for all k . For the inductive step, suppose $n > 1$. Let T^* be a conservative decision tree of minimum expected cost for the instance. Assume for contradiction that strategy $T_k^c(\pi)$ does not have minimum expected cost. To simplify the proof, we assume in what follows that the values $c_i/(1 - p_i)$ are distinct for all tests i . It is not difficult to eliminate this assumption. Index the tests so that $c_1/(1 - p_1) < c_2/(1 - p_2) < \dots < c_n/(1 - p_n)$.

Tree T^* has a root labeled with a test i , and a left and right subtree T_L^* and T_R^* . After seeing the value of test i at the root of T^* , we still need to look for either $k - 1$ or k additional 0's. Thus T_L^* and T_R^* are min-cost solutions for induced $(k - 1)$ -of- $(n - 1)$ and k -of- $(n - 1)$ testing instances respectively,

Consider running the algorithm on both of these induced instances. The algorithm outputs the same permutation $\tilde{\pi} = (1, 2, \dots, i - 1, i + 1, \dots, n)$ in both cases. By induction, $T_{k-1}^c(\tilde{\pi})$ and $T_k^c(\tilde{\pi})$ are optimal for

the induced instances. Let T' be the tree produced from T^* by replacing the optimal left and right subtrees of T^* by the optimal subtrees $T_{k-1}^c(\tilde{\pi})$ and $T_k^c(\tilde{\pi})$ respectively. Clearly, T' is also an optimal tree for the original problem instance.

Let $\pi' = (i, 1, \dots, i-1, i+1, \dots, n)$. Thus $T' = T_k^c(\pi')$. If $i = 1$, then $\pi = \pi'$, and because $T_k^c(\pi')$ is optimal, so is $T_k^c(\pi)$. This contradicts our assumption that $T_k^c(\pi)$ is not optimal. Therefore, $i \neq 1$.

Now consider permutation $\pi' = (i, 1, 2, \dots, i-1, i+1, \dots, n)$. Move i forward in the permutation, one spot at a time, until the resulting permutation π'' is such that $T_k^c(\pi'')$ is not optimal. Since $T_k^c(\pi)$ is not optimal, either $\pi'' = 1, 2, \dots, j-1, i, j, \dots, i-1, i+1, \dots, n$ for some $j < i$, or $\pi'' = 1, 2, \dots, i-1, i, i+1, \dots, n = \pi$ (in which case, let $j = i$).

By the definition of j , swapping $j-1$ and i in π'' yields a permutation $\hat{\pi}$ whose associated strategy $T_k(\hat{\pi})$ is optimal. We now compare the expected costs of $T_k^c(\pi'')$ and $T_k(\hat{\pi})$ on an element $x \in \{0, 1\}^n$. If the values of tests $1, \dots, j-2$ on input x include at least k 0's, or strictly fewer than $k-1$ 0's, then the cost of using $T_k^c(\pi'')$ and $T_k(\hat{\pi})$ on x are equal. If the values of tests $1, \dots, j-2$ on input x include exactly $k-1$ 0's, then the costs of $T_k^c(\pi'')$ and $T_k(\hat{\pi})$ may differ. Both trees perform tests $1, \dots, j-2$ first. If $x_{j-1} = 0$, then $T_k^c(\pi'')$ will perform only that one additional test, while if $x_{j-1} = 1$, then $T_k^c(\pi'')$ will perform both tests $j-1$ and i . Similarly, if $x_i = 0$, then $T_{\hat{\pi}}$ will perform only the one additional test i , while if $x_i = 1$, then $T_{\hat{\pi}}$ will perform both tests $j-1$ and i . Any subsequent testing will be the same in both trees. Because the expected cost of $T_k^c(\pi'')$ is greater than the expected cost of $T_k(\hat{\pi})$, $c_{j-1} + p_{j-1}c_i > c_i + p_i c_{j-1}$, and thus $c_{j-1}/(1 - p_{j-1}) > c_i/(1 - p_i)$.

Since $j \leq i$, this violates the sorted ordering of π , which is a contradiction. \square

5 The Algorithm for the CMT(k) problem

We begin with some useful lemmas. The algorithms of Condon et al. for maximizing throughput of 1-of- n testing rely crucially on the fact that saturation of all processors implies optimality. We show that the same holds for conservative k -of- n testing.

Lemma 2. *Let R be a routing for an instance of the conservative k -of- n testing problem. If R saturates all processors, then it is optimal.*

Proof. Each processor O_i can test at most r_i items per unit time. Thus at processor O_i , there are at most $r_i(1 - p_i)$ tests performed that have the value 0. Let f denote the k -of- n function.

Suppose R is a routing achieving throughput F . Since F items enter the system per unit time, F items must also leave the system. An item x such that $f(x) = 0$ does not leave the system until it fails k tests. An item x such that $f(x) = 1$ does not leave the system until it has had all tests performed on it. Thus, per unit time, in the entire system, the number of tests performed that have the value 0 must be $F \times M$, where $M = (k \cdot \text{Prob}[x \text{ has at least } k \text{ 0's}] + \sum_{j=0}^{k-1} j \cdot \text{Prob}[x \text{ has exactly } j \text{ 0's}])$.

Since at most $r_i(1 - p_i)$ tests with the value 0 can occur per unit time at processor O_i , $F \times M \leq \sum_{i=1}^n r_i(1 - p_i)$. Solving for F , this gives an upper bound of $F \leq \sum_{i=1}^n r_i(1 - p_i)/M$ on the maximum throughput. This bound is tight if all processors are saturated, and hence a routing saturating all processors achieves the maximum throughput. \square

In the above proof, we rely on the fact that every routing with throughput F results in the same number of 0 test values being generated in the system per unit time. Note that this is not the case for *standard* testing, where the number of 0 test values generated can depend on the routing itself, and not just on the throughput of that routing. We now give a simple counterexample showing that, in fact, saturation does not imply optimality for the SMT(k) problem. Consider the MAXTHROUGHPUT 2-of-3 testing instance where $p_1 = 1/2, p_2 = 1/4, p_3 = 3/4$, and $r_1 = 2, r_2 = 1\frac{3}{4}, r_3 = 1\frac{3}{4}$.

The following is a 2-of-3 testing strategy: *Given item x , perform test 1. If $x_1 = 1$, follow strategy $T_1^s(\pi_1)$, where $\pi_1 = (2, 3)$. Else if $x_1 = 0$, follow strategy $T_1^s(\pi_2)$, where $\pi_1 = (3, 2)$.*

Assigning 2 units of flow to this strategy saturates the processors: O_1 is saturated since it receives the 2 units entering the system, O_2 is saturated since it receives $1 = 2 \times p_1$ units from O_1 and $3/4 = 2 \times p_3 \times (1 - p_1)$

items from O_1, O_2 . Similarly, O_3 is saturated since it receives $1 = 2 \times (1 - p_1)$ units from O_1 and $3/4 = 2 \times (1 - p_3) \times p_1$ units from $O_1 O_3$.

We show that the routing is not optimal by giving a different routing with higher throughput. The routing uses two strategies. The first is as follows: *Given item x , perform test 1. If $x_1 = 1$, follow strategy $T_2^s(\pi_1)$, where $\pi_1 = (3, 2)$. Else, if $x_1 = 0$ follow strategy $T_1^s(\pi_1)$, where $\pi_2 = (2, 3)$.* The second strategy used by the routing is $T_2^s(\pi_3)$, where $\pi_3 = (3, 2, 1)$. Assigning $F = 1\frac{1}{2}$ units to the first strategy uses $1\frac{1}{2}$ units of the capacity of O_1 , $15/16 = 1\frac{1}{2} \times (1 - p_1) + 1\frac{1}{2} \times p_1 \times (1 - p_3)$ units of the capacity of O_2 , and $15/16 = 1\frac{1}{2} \times (1 - p_1) + 1\frac{1}{2} \times (1 - p_1) \times p_2$ of the capacity of O_3 . This leaves O_2 and O_3 with residual capacity more than $3/4 < 1 + 3/4 - 15/16$, and O_1 with residual capacity $1/2 = 2 - 1\frac{1}{2}$. We can then assign $3/4$ additional units to the second strategy without violating any of the rate constraints, for a routing with total throughput $2\frac{1}{4}$. (The resulting routing is not optimal, but illustrates our point.)

The routing produced by our algorithm for the CMT(k) problem uses only strategies of the form $T_k^c(\pi)$, for some permutation π of the tests (in terms of the LP, this means $z_T > 0$ only if $T = T_k^c(\pi)$ for some π). We call such a routing a *permutation routing*. We say that it has a *saturated suffix* if for some subset Q of the processors (1) R saturates all processors in Q , and (2) for every strategy $T_k^c(\pi)$ used by R , the processors in Q appear together in a suffix of π .

With this definition, and the above lemma, we are now able to generalize a key lemma of Condon et al. to apply to conservative k -of- n testing. The proof is essentially the same as theirs; we present it below for completeness.

Lemma 3. (Saturated Suffix Lemma) *Let R be a permutation routing for an instance of the CMT(k) problem. If R has a saturated suffix, then R is an optimal solution for the instance.*

Proof. If R saturates all processors, then the previous lemma guarantees its optimality. If not, let L denote the set of processors not saturated by R . Imagine that we removed the rate constraints for each processor in L . Let R' be an optimal routing for the resulting problem. We may assume that on any input x , R' performs the tests in L in some fixed arbitrary order (until and unless k tests with value 0 are obtained), prior to performing any tests in Q . This assumption is without loss of generality, because if not, we could modify R' to first perform the tests in L without violating feasibility, since the processors in L have no rate constraints, and performing their tests first can only decrease the load on the other processors. Thus the throughput attained by R' is $T_R \times \frac{1}{p_L}$, where T_R denotes the maximum throughput achievable just with the operators in Q , and p_L is the probability that a random x will have the value 0 for fewer than k of the tests in L (i.e. it will not be eliminated by the tests in L).

Routing R also routes flow first through L , and then through Q . Since it saturates the operators in Q , by the previous lemma, it achieves maximum possible throughput with those operators. It follows that R achieves the same throughput as R' , and hence is optimal for the modified instance where processors in L have no rate constraints. Since removing constraints can only increase the maximum possible throughput, it follows that R is also optimal for the original instance. \square

5.1 The Equal Rates Case

We begin by considering the CMT(k) problem in the special case where the rate limits r_i are equal to some constant value r for all processors. Condon et al. presented a closed-form solution for this case when $k = 1$ [4]. The solution is a permutation routing that uses n strategies of the form $T_1(\pi)$. Each permutation π is one of the n left cyclic shifts of the permutation $(1, \dots, n)$. More specifically, for $i \in \{1, \dots, n\}$, let $\pi_i = (i, i+1, \dots, n, 1, 2, \dots, i-1)$, and let $T_i = T_1^c(\pi_i)$. The solution assigns $r(1 - p_{i-1})/(1 - p_1 \cdots p_n)$ units of flow to each T_i . By simple algebra, Condon et al. verified that the solution saturates all processors. Hence it is optimal.

The solution of Condon et al. is based on the fact that for the 1-of- n problem, assigning $(1 - p_{i-1})$ flow to each T_i equalizes the load on the processors. Surprisingly, this same assignment equalizes the load for the k -of- n problem as well. Using this fact, we obtain a closed-form solution to the CMT(k) problem.

Lemma 4. Consider an instance of the CMT(k) problem where all processors have the same rate limit r . For $i \in \{1, \dots, n\}$, let T_i be as defined above. Let $X_{a,b} = \sum_{\ell=a}^b (1 - x_\ell)$. The routing that assigns $r(1 - p_{i-1})/\alpha$ items per unit time to strategy T_i saturates all processors, where $\alpha = \sum_{t=1}^k \text{Prob}[X_{1,n} \geq t]$.

Proof. We begin by considering the routing in which $(1 - p_{i-1})$ units of flow are assigned to each T_i . Consider the question of how much flow arrives per unit time at processor 1, under this routing. For simplicity, assume now that $k = 2$. Thus as soon as an item has failed 2 tests, it is discarded. Let $q_i = (1 - p_i)$.

Of the q_n units assigned to strategy T_1 , all q_n arrive at processor 1. Of the q_{n-1} units assigned to strategy T_n , all q_{n-1} arrive at processor 1, since they can fail either 0 or 1 test (namely test n) beforehand.

Of the q_{n-2} units assigned to strategy T_{n-1} , the number reaching processor O_1 is $q_{n-2}\beta_{n-1}$, where β_{n-1} is the probability that an item fails either 0 or 1 of tests $n-1$ and n . Therefore, $\beta_{n-1} = 1 - q_{n-1}q_n$. More generally, for $i \in \{1, \dots, n\}$, of the q_{i-1} units assigned to T_i , the number reaching processor 1 is $q_{i-1}\beta_i$, where β_i is the probability that a random item fails a total of 0 or 1 of tests $i, i+1, \dots, n$. Thus, $\beta_i = \text{Prob}[X_{i,n} = 0] + \text{Prob}[X_{i,n} = 1]$. It follows that the total flow arriving at Processor 1 is $\sum_{i=1}^n q_{i-1}(\text{Prob}[X_{i,n} = 0]) + \sum_{i=1}^n q_{i-1}(\text{Prob}[X_{i,n} = 1])$.

Consider the second summation, $\sum_{i=1}^n (q_{i-1}\text{Prob}[X_{i,n} = 1])$. We claim that this summation is equal to $\text{Prob}[X_{1,n} \geq 2]$, which is the probability that x has *at least* two x_i 's that are 0. To see this, consider a process where we observe the value of x_n , then the value of x_{n-1} and so on down towards x_1 , stopping if and when we have observed exactly two 0's. The probability that we will stop at some point, having observed two 0's, is clearly equal to the probability that x has *at least* two x_i 's that are set to 0. The condition $\sum_{j=i}^n (1 - x_j) = 1$ is satisfied when exactly 1 of x_n, x_{n-1}, \dots, x_i has the value 0. Thus $q_{i-1}(\text{Prob}[X_{i,n} = 1])$ is the probability that we observe exactly one 0 in x_n, \dots, x_i , and then we observe a second 0 at x_{i-1} . That is, it is the probability that we stop after observing x_{i-1} . Since the second summation takes the sum of $q_{i-1}\text{Prob}[X_{i,n} = 1]$ over all i between 1 and n , the summation is precisely equal to the probability of stopping at some point in the above process, having seen two 0's. This proves the claim.

An analogous argument shows that the first summation, $\sum_{i=1}^n (q_{i-1}\text{Prob}[X_{i,n} = 0])$, is equal to $\text{Prob}[X_{1,n} \geq 1]$.

It follows that the amount of flow reaching Processor 1 is $\text{Prob}[X_{1,n} \geq 1] + \text{Prob}[X_{1,n} \geq 2]$. This expression is symmetric in the processor numbers, so the amount of flow reaching every O_i is equal to this value. Thus the above routing causes all processors to receive the same amount of flow. Hence, if all processors have the same rate limit r , scaling this routing by an appropriate multiplicative factor will yield a routing that saturates all processors. More particularly, the routing that assigns $r q_{i-1}/(\text{Prob}[X_{1,n} \geq 1] + \text{Prob}[X_{1,n} \geq 2])$ units to each strategy T_i will saturate all processors if they have common rate limit r .

The above argument for $k = 2$ can easily be extended to arbitrary k . The resulting saturating routing for arbitrary k , when all processors have rate limit r , assigns $r q_{i-1}/(\sum_{t=1}^k \text{Prob}[X_{1,n} \geq t])$ items per unit time to strategy T_i . \square

5.2 The Equalizing Algorithm of Condon et al.

Our algorithm for the CMT(k) problem is an adaptation of one of the two MAXTHROUGHPUT algorithms, for the special case where $k = 1$, given by Condon et al. [4]. We begin by reviewing that algorithm, which we will call the *Equalizing Algorithm*. Note that when $k = 1$, it only makes sense to consider strategies that are permutation routings, since an item can be discarded as soon as it fails a single test.

Consider the CMT(k) problem for $k = 1$. View the problem as one of constructing a flow of items through the processors. The capacity of each processor is its rate limit, and the amount of flow sent along a permutation π (i.e., assigned to strategy $T_1^c(\pi)$) is equal to the number of items sent along that path per unit time. Sort the tests by their rate limits, and re-number them so that $r_n \geq r_{n-1} \geq \dots \geq r_1$. Assume for the moment that all rate limits r_i are distinct.

The Equalizing Algorithm constructs a flow incrementally as follows. Imagine pushing flow along the single permutation $(n, \dots, 1)$. Suppose we continuously increase the amount of flow being pushed, beginning from zero, while monitoring the “residual capacity” of each processor, i.e., the difference between its rate

limit and the amount of flow it is already receiving. (For the moment, let us not worry about exceeding the rate limit of a processor.)

Consider two adjacent processors, i and $i - 1$. As we increase the amount of flow, the residual capacity of each decreases continuously. Initially, at zero flow, the residual capacity of i is greater than the residual capacity of $i - 1$. It follows by continuity that the residual capacity of i cannot become less than the residual capacity of $i - 1$ without the two residual capacities first becoming equal. We now impose the following stopping condition: increase the flow sent along permutation $(n, \dots, 1)$ until either (1) some processor becomes saturated, or (2) the residual capacities of at least two of the processors become equal. The second stopping condition ensures that when the flow increase is halted, permutation $(n, \dots, 1)$ still orders the processors in decreasing order of their residual capacities. (Algorithmically, we do not increase the flow continuously, but instead directly calculate the amount of flow which triggers the stopping condition.)

If stopping condition (1) above holds when the flow increase is stopped, then the routing can be shown to have a saturated suffix, and hence it is optimal.

If stopping condition (2) holds, we keep the current flow, and then augment it by solving a new MAX-THROUGHPUT problem in which we set the rate limits of the processors to be equal to their residual capacities under the current flow (their p_i 's remain the same).

We solve the new MAXTHROUGHPUT problem as follows. We group the processors into equivalence classes according to their rate limits. We then replace each equivalence class with a single mega-processor, with a rate limit equal to the rate limit of the constituent processors, and probability p_i equal to the product of their probabilities. We then essentially apply the procedure for the case of distinct rate limits to the mega processors.

The one twist is the way in which we translate flow sent through a mega-processor into flow sent through the constituent processors of that mega-processor; we route the flow through the constituent processors so as to equalize their load. We accomplish this by dividing the flow proportionally between the cyclic shifts of a permutation of the processors, using the same proportional allocation as used in the saturating routing of Lemma 4. We thus ensure that the processors in each equivalence class continue to have equal residual capacity. Note that, under this scheme, the residual capacity of a processor in a mega-processor may decrease more slowly than it would if all flow were sent directly to that processor (because some flow may first be filtered through other processors in the mega-processor) and this needs to be taken into account in determining when the stopping condition is reached.

Here is an example illustrating the Equalizing Algorithm. We can observe how our algorithm works on the following 1-of-3 CMT(k) problem (which, since $k = 1$ is the same as SMT(k) problem.) Suppose we have 3 processors, O_1, O_2, O_3 with rate limits $r_1 = 3, r_2 = 14$, and $r_3 = 18$, and probabilities $p_1 = 1/8, p_2 = 1/2$ and $p_3 = 1/3$. When flow is sent along O_3, O_2, O_1 , after 6 units of flow is sent we have a stopping condition when O_3 and O_2 , have the same residual capacity of 12; the residual capacity of O_1 is 2.

Our algorithm then does a recursively call where the operators O_3 and O_2 are combined into a mega-processor $O_{2,3}$; which has $p_{2,3} = 1/2 \times 1/3 = 1/6$. How flow is sent through the mega-processor, $O_{2,3}$ is by sending $3/7$ fraction through O_3, O_2 and $4/7$ fraction through O_2, O_3 ; we observe that for one unit of flow sent through $O_{2,3}$ the amount of capacity used by each processor is $3/7 + 2/7 = 5/7$. Flow is now sent along $O_{2,3}, O_1$; after 12 units of flow, we reach a stopping condition when O_1 is saturated. Even though O_2 and O_3 are not saturated (they have $12 - 12 \times 5/7$ residual capacity left) the flows constructed as described provide optimal throughput.

The Equalizing Algorithm, implemented in a straightforward way, produces a routing that may use an exponential number of different permutations. Condon et al. describe methods for reducing the number of permutations used [4].

5.3 An Equalizing Algorithm for the CMT(k) problem

We prove the following theorem.

Theorem 5. *There is a $O(kn^2)$ combinatorial algorithm for solving the CMT(k) problem.*

Proof. We extend the Equalizing Algorithm of Condon et al., to apply to arbitrary values of k .

Again, we will push flow along the permutation of the processors $(n, \dots, 1)$ (where $r_n \geq r_{n-1} \geq \dots \geq r_1$) until one of the two stopping conditions is reached: (1) a processor is saturated, or (2) two processors have equal residual capacity. Here, however, we do not discard an item until it has failed k tests, rather than discarding it as soon as it fails one test. To reflect this, we divide the flow into k different types, numbered 0 through $k-1$, depending on how many tests its component items have failed. Flow entering the system is all of type 0.

When m flow of type τ enters a processor O_i , $p_i m$ units pass test O_i , and $(1 - p_i)m$ units fail it. So, if $\tau < k-1$, then of the m incoming units of type τ , $(1 - p_i)m$ units will exit processor O_i as type $\tau+1$ flow, and $p_i m$ will exit as type τ flow. Both types will be passed on to the next processor in the permutation, if any. If $\tau = k-1$, then $p_i m$ units will exit as type τ flow and be passed on to the next processor, and the remaining $(1 - p_i)m$ will be discarded.

Algorithmically, we need to calculate the minimum amount of flow that triggers a stopping condition. This computation is only slightly more complicated for general k than it is for $k=1$. The key is to compute, for each processor O_i , what fraction of the flow that is pushed into the permutation will actually reach processor O_i (i.e. we need to compute the quantity $g(T_k^c(\pi), i)$ in the LP.)

If stopping condition (2) holds, we keep the current flow, and augment it by solving a new MAXTHROUGHPUT problem in which we set the rate limits of the processors to be equal to their residual capacities under the current flow (their p_i 's remain the same.) To solve the new MAXTHROUGHPUT problem, we again group the processors into equivalence classes according to their rate limits, and replace each equivalence class with a single mega-processor, with a rate limit equal to the rate limit of the constituent processors, and probability p_i equal to the product of their probabilities.

We then want to apply the procedure for the case of distinct rate limits to the mega processors. To do this, we need to translate flow sent into a mega-processor into flow sent through the constituent processors of that mega-processor, so as to equalize their load. We do this translation separately for each type of flow entering the mega-processor. Flow of type τ entering the mega-processor is sent into the constituent processors of the mega-processor according to the closed-form solution for $(k-\tau)$ -of- n' testing, where n' is the number of constituent processors of the mega-processor, because it will be discarded if it fails $k-\tau$ more tests. We also need to compute how much flow of each type ends up leaving the mega-processor (some of the incoming flow of type τ entering the mega-processor may, for example, become outgoing flow of type $\tau+n'$), and how much its residual capacity is reduced by the incoming flow. The algorithm can be implemented to run in time $O(kn^2)$. We give further details, with pseudocode below. \square

Pseudocode The pseudocode is presented below. The following information will be helpful in understanding it.

At each stage of the algorithm, the processors are partitioned into equivalence classes. The processors in each equivalence class constitute a mega-processor. Each equivalence class consists of a contiguous subsequence of processors, in the sorted sequence O_n, O_2, \dots, O_1 . We use m to denote the number of mega-processors (equivalence classes). The processors in each equivalence class all have the same residual capacity. In Step 1 of the algorithm, we partition the processors according to which have the same residual capacity. We use O_l to denote the mega-processor containing the processors in equivalence class E_l .

In Step 2, we compute the amount of flow \hat{t} that triggers one of the two stopping conditions. In order to do this, we need to know that rate at which the residual capacity of each processor within an equivalence class E_i will be reduced when flow is sent down the mega-processors in the order E_m, \dots, E_1 . We use $\xi(i)$ to denote the amount by which the residual capacity of the processors in E_i is reduced when one unit of flow is sent in that order.

The equation for $\xi(i)$ follows from the lemmas and proof for the algorithm. We use $f_j(z)$ to denote the amount of flow that would reach processor z , if one unit of flow were sent down the permutation O_n, \dots, O_1 , where these are the original processors, not the mega-processors. This is precisely equal to the probability that random item x has fewer than k 0's in tests $n, \dots, z+1$. We compute the value of $f_j(z)$ for all z and k in a separate initialization routine. The key here is noticing that if you send one unit of flow down the

megaprocessors E_m, \dots, E_1 , the amount of flow reaching mega-processor i is precisely $f_j(c(i))$, where $c(i)$ is the highest index of a processor in E_i ; the amount of flow reaching the megaprocessor depends only on how many 0's have been encountered in test $n, \dots, c(i) + 1$, on the order used to perform those tests.

The quantity \hat{t}_1 is the amount of flow sent down E_m, \dots, E_1 that would cause saturation of the processors in E_1 . The quantity \hat{t}_2 is the minimum amount of flow sent down E_m, \dots, E_1 that would cause the residual capacities of two megaprocessors to equalize. The stopping condition holds at the minimum of these two quantities.

The algorithm outputs a routing as a list of pairs of the form $((E_m, \dots, E_1), \hat{t})$, meaning that \hat{t} flow should be sent down the permutation of megaprocessors (E_m, \dots, E_1) . Of course, flow coming into each mega-processor is routed so as to equalize the load on each of its constituent processors.

It is easy to see that the algorithm makes at most n recursive calls, because mega-processors can only be merged a total of $n - 1$ times. Excluding the computation of $\xi(i)$, the time spent in each recursive call is clearly $O(n)$. During each recursive call, the value of $\xi(i)$ can be computed in time $\mathcal{O}(nk)$ via dynamic programming. This yields a total running time of $\mathcal{O}(n^2k)$.

Algorithm: MAXTHROUGHPUT Initialization

```

 $f_i(j) \leftarrow 0, \forall j \in \{1, \dots, n\}, \forall i \in \{0, \dots, k - 1\};$ 
 $f_0(1) \leftarrow 1;$ 
for  $(j \leftarrow 2, i \leq n; i \leftarrow i + 1)$  do
  for  $(i \leftarrow 0, j \leq k - 1; j \leftarrow j + 1)$  do
     $f_i(j) \leftarrow q_{j-1}f_{i-1}(j-1) + p_{j-1}f_i(j-1);$ 
return SolveMaxThroughput( $p_1, \dots, p_n, r_1, \dots, r_n$ );
```

Example We illustrate our algorithm for the CMT(k) problem on the following example, where $k = 2$ and $n = 4$. Suppose the probabilities are $p_1 = p_2 = p_3 = 1/2$, $p_4 = 3/4$, and the rate limits are $r_1 = r_2 = 12$, $r_3 = r_4 = 10$.

We will use the following fact, which is an easy corollary of Lemma 4. The strategies T_i are the cyclic permutation strategies defined in that Lemma. It follows immediately from what is shown in the proof of the lemma, namely that assigning q_{i-1} flow unit to each T_i equalizes the load on each processor.

Fact: Given processors O_1, \dots, O_n , if R is a routing that assigns a $\frac{q_{i-1}}{\sum_{j=1}^n q_j}$ fraction of the total flow to strategy T_i , then R uses the same amount of capacity in each processor.

Our algorithm first combines processors with same rate limits into mega-processors; thus we combine O_1 and O_2 into mega-processor $O_{1,2}$ with rate limit 12. It routes flow through this mega-processor by sending 1/2 fraction of the flow in the order O_1, O_2 , and sending the other 1/2 fraction in the order O_2, O_1 . Similarly, O_3 and O_4 have the same rate limit, so they are combined into a mega-processor $O_{3,4}$ with rate limit 10, where 1/3 fraction of the flow is sent along O_3, O_4 , and 2/3 fraction of the flow is sent along O_4, O_3 .

Our mega-processor $O_{1,2}$ has a higher rate limit than $O_{3,4}$, consequently our algorithm routes flow in the order $O_{1,2}, O_{3,4}$. We now show that the stopping condition is reached after sending 6 units of flow along this route, since after the 6 units of flow have been sent we have equalized the residual capacity of all the operators without saturating any of them.

The 6 units of flow decreased the capacity of processors O_1 , and O_2 in $O_{1,2}$ by 6, since $k = 2$ and thus flow cannot be discarded before it has been subject to at least two tests.

We now calculate the reduction of capacity in O_3 and O_4 caused by the 6 unit of flow sent through $O_{1,2}, O_{3,4}$. Flow leaving $O_{1,2}$ has a 1/4 probability of having failed both processors in $O_{1,2}$ and exiting the system; for flow that stays in the system to be tested by $O_{3,4}$, it has a 1/4 chance of having passed the test of both processors; it has a 1/2 chance of having passed the test of one processor and having failed the test of the other processor. Thus, of the 6 units of flow sent into $O_{1,2}$, $1/4 \times 6 = 3/2$ units are passed on to $O_{3,4}$ as type 0 flow, and $1/2 \times 6 = 3$ units of flow are passed on to $O_{3,4}$ as type 1 flow.

Of the $3/2$ units of type 0 flow, entering $O_{3,4}$, all of it must undergo both test 3 and test 4, since flow is

```
SolveMaxThroughput( $p_1, \dots, p_n, r_1, \dots, r_n$ )
```

Input: n selectivities p_1, \dots, p_n ; n rate limits $r_1 \leq \dots \leq r_n$

Output: compact representation of solution to the MAXTHROUGHPUT problem for the given input parameters

- 1.** // form the equivalence classes E_m, \dots, E_1 ;
 $E_1 \leftarrow \{O_1\}$;
 $m \leftarrow 1$; // m is the number of equivalence classes
 $R_1 \leftarrow r_1$;

for $(\ell \leftarrow 2; \ell \leq n; \ell \leftarrow \ell + 1)$ **do**
 if $(r_\ell \neq r_{\ell-1})$ **then**
 $m \leftarrow m + 1$;
 $E_m \leftarrow \{O_\ell\}$;
 $R_m \leftarrow r_\ell$;
 else // $(r_\ell == r_{\ell-1})$
 $E_m \leftarrow E_m \cup \{O_\ell\}$;
- 2.** // calculate \hat{t} using the following steps;
for $(i \leftarrow 1; i \leq m; i \leftarrow i + 1)$ **do**
 $c(i) \leftarrow$ highest index of an operator in E_i ;
 $b(i) \leftarrow$ lowest index of an operator in E_i ;
 $\xi(i) \leftarrow \sum_{j=0}^{k-1} f_j(c(i)) \cdot \left(\sum_{v=1}^{k-j} \Pr \left[\sum_{t=b(i)}^{c(i)} (1 - x_t) \geq v \right] \right)$;
 $\hat{t}_1 \leftarrow \frac{R_1}{\xi(1)}$;
 $\hat{t}_2 \leftarrow \min_{i \in [2, \dots, m]} \left(\frac{R_i - R_{i-1}}{\xi(i) - \xi(i-1)} \right)$;
 $\hat{t} \leftarrow \min(\hat{t}_1, \hat{t}_2)$;
- 3.** // calculate the residual capacity for each operator O_ℓ ;
for $(\ell \leftarrow 1; \ell \leq n; \ell \leftarrow \ell + 1)$ **do**
 $j \leftarrow$ index of the equivalence class E_j containing operator O_ℓ ;
 $r'_\ell \leftarrow r_\ell - \xi(j)\hat{t}$;
- 4.** // $K = (\{E_m, \dots, E_1\}, \hat{t})$;
if $(r'_1 == 0)$ **then** // residual capacity of equivalence class E_1 is 0
 return K ;
else
 $K' \leftarrow \text{SolveMaxThroughput}(p_1, \dots, p_n, r'_1, \dots, r'_n)$;
 return $K \circ K'$; // i.e. the concatenation of K and K'

not discarded until it has failed two tests. Thus that flow reduces the capacity of both O_3 and O_4 by $3/2$ units.

Of the 3 units of type 1 flow entering $O_{3,4}$, $1/3$ is tested first by O_3 , and then by O_4 only if it passes test 3 (which it does with probability $1/2$). The remaining $2/3$ is tested first by O_4 , and then by O_3 only if it passes test 4 (which is does with probability $3/4$). Thus of the 3 units of type 1 flow, $3 \times (1/3 + 2/3 \times 3/4) = 5/2$ units reach O_3 , and $3 \times (2/3 + 1/3 \times 1/2) = 5/2$ units reach O_4 .

Hence the $3+3/2$ total units of flow entering $O_{3,4}$ reduce the capacities of both O_3 and O_4 by $5/2+3/2 = 4$.

We have thus shown that the 6 units of flow sent first to $O_{1,2}$ and then to $O_{3,4}$, cause the residual capacities of O_1 and O_2 to be $12 - 6 = 6$, and the residual capacities of O_3 and O_4 to be $10 - 4 = 6$. Thus the residual capacities of all operators equalize, as claimed.

At this point our algorithm constructs a new mega-processor, by combining the processors in $O_{1,2}$ with the processors in $O_{3,4}$. All the processors in the resulting mega-processor, $O_{1,2,3,4}$, have a residual capacity of 6. Using the proportional allocation in the routing of Lemma 4 to route flow into $O_{1,2,3,4}$, we will route $1/7$ fraction of the flow along $\pi_1 = \{1, 2, 3, 4\}$, $2/7$ fraction of the flow along the route $\pi_2 = \{2, 3, 4, 1\}$, $2/7$ fraction of the flow along the route $\pi_3 = \{3, 4, 1, 2\}$, and $2/7$ fraction of flow along the route $\pi_4 = \{4, 1, 2, 3\}$.

By sending 7 units of flow through $O_{1,2,3,4}$ we decrease each processor's residual capacity by 6; thus saturating all processors.

Our final routing achieves a throughput of $6 + 7 = 13$ which is optimal.

6 An Ellipsoid-Based Algorithm for the SMT(k) problem

There is a simple and elegant algorithm that solves the MINCOST problem for standard k -of- n testing, due to Salloum, Breuer, and (independently) Ben-Dov [9, 10, 1]. It outputs a strategy compactly represented by two permutations, one ordering the operators in increasing order of the ratio $c_i/(1 - p_i)$, and the other in increasing order of the ratio c_i/p_i . Chang et al. and Salloum and Breuer later gave modified versions of this algorithm that output a less compact, but more efficiently evaluable representation of the same strategy [11, 2].

We now show how to combine previous techniques to obtain a polynomial-time algorithm for the SMT(k) problem based on the ellipsoid method. The algorithm uses a technique of Deshpande and Hellerstein [5]. They showed that, for 1-of- n testing, an algorithm solving the MINCOST problem can be combined with the ellipsoid method to yield an algorithm for the MAXTHROUGHPUT problem. In fact, as we see in the proof below, their approach is actually a generic one, and can be applied to testing of other functions.

The ellipsoid-based algorithm for k -of- n testing makes use of the dual of the LP for the CMT(k) problem, which is as follows:

Dual of Max-Throughput LP: Given $r_1, \dots, r_n > 0$, $p_1, \dots, p_n \in (0, 1)$, find an assignment to the variables y_i , for all $i \in \{1, \dots, n\}$, minimizing

$$F = \sum_{i=1}^n r_i y_i$$

subject to the constraints:

- (1) $\sum_{i=1}^n g(\pi, i) y_i \geq 1$ for all $T \in \mathcal{T}_c$ and
- (2) $y_i \geq 0$ for all $i \in \{1, \dots, n\}$.

Theorem 6. *There is a polynomial-time algorithm, based on the ellipsoid method, for solving the SMT(k) problem.*

Proof. The approach of Deshpande and Hellerstein works as follows. The input consists of the p_i and the r_i , and the goal is to solve the MAXTHROUGHPUT LP in time polynomial in n . The number of variables of the MAXTHROUGHPUT LP is not polynomial, so the LP cannot be solved directly. Instead, the idea is

to solve it by first using the ellipsoid method to solve the dual LP. The ellipsoid method is run using an algorithm that simulates a separation oracle for the dual in time polynomial in n . During the running of the ellipsoid method, the violated constraints returned by the separation oracle are saved in a set M . Each constraint of the dual corresponds to an ordering T . When the ellipsoid method terminates, a modified version of the MAXTHROUGHPUT LP is generated, which includes only the variables z_T corresponding to orderings T in M (i.e. the other variables z_T are set to 0). This modified version can then be solved directly using a polynomial-time LP algorithm. The resulting solution is an optimal solution for the original MAXTHROUGHPUT LP.

The above approach requires a polynomial-time algorithm for simulating the separation oracle for the dual. Deshpande and Hellerstein's method for simulating the separation oracle relies on the following observations. In the dual LP for the MAXTHROUGHPUT 1-of- n testing problem, there are $n!$ constraints corresponding to the $n!$ permutations of the processors. The constraint for permutation π is $\sum_{i=1}^n g(T_1(\pi), i)y_i \leq 1$. If one views y as a vector of costs, where the cost of i is y_i , then $\sum_{i=1}^n g(T, i)y_i$ is the expected cost of testing an item x using ordering T . Thus one can determine the ordering T that minimizes $\sum_{i=1}^n g(T, i)y_i$ by solving the MINCOST problem with probabilities p_1, \dots, p_n and cost vector y . (Liu et al.'s approximation algorithm for generic MAXTHROUGHPUT also relies on that observation [8].)

If the MINCOST ordering T has expected cost less than 1, then the constraint it corresponds to is violated. Otherwise, since the right hand side of each constraint is 1, y obeys all constraints. Thus simulating the separation oracle for the dual on input y can be done by first running the MINCOST algorithm (with probabilities p_i and costs y_i) to find a MINCOST ordering T . Once T is found, the values of the coefficients $g(T, i)$ are calculated. These are used to calculate $\sum_{i=1}^n g(T, i)$, the expected cost of T . If this value is less than 1, then the constraint $\sum_{i=1}^n g(T, i)$ is returned.

To apply the above approach to MAXTHROUGHPUT for standard k -of- n testing, we observe that in the dual LP for this problem, there is a constraint, $\sum_{i=1}^n g(T, i)y_i \leq 1$, for every possible strategy T . We can simulate a separation oracle for the dual on input y by running a MINCOST algorithm for standard k -of- n testing. We also need to be able to compute the $g(T, i)$ values for the strategy output by that algorithm. The algorithm of Chang et al. MINCOST standard k -of- n testing problem is suitable for this purpose, as it can easily be modified to output the $g(T, i)$ values associated with its output strategy T [2]. \square

References

- [1] Y. Ben-Dov. Optimal testing procedure for special structures of coherent systems. *Management Science*, 27:1410–1420, 1981.
- [2] Ming-Feng Chang, Weiping Shi, and W. Kent Fuchs. Optimal diagnosis procedures for k-out-of-n structures. *IEEE Trans. Computers*, 39(4):559–564, 1990.
- [3] Anne Condon, Amol Deshpande, Lisa Hellerstein, and Ning Wu. Flow algorithms for two pipelined filter ordering problems. In *PODS*, pages 193–202. ACM, 2006.
- [4] Anne Condon, Amol Deshpande, Lisa Hellerstein, and Ning Wu. Algorithms for distributional and adversarial pipelined filter ordering problems. *ACM Transactions on Algorithms*, 5(2), 2009.
- [5] Amol Deshpande and Lisa Hellerstein. Parallel pipelined filter ordering with precedence constraints. Pre-publication version available at <http://cis.poly.edu/~hstein/pubs/filterwithconstraint.pdf>, To appear in ACM Transactions on Algorithms.
- [6] M. R. Garey. Optimal task sequencing with precedence constraints. *Discrete Math.*, 4:37–56, 1973.
- [7] Murali S. Kodialam. The throughput of sequential testing. In Karen Aardal and Bert Gerards, editors, *IPCO*, volume 2081 of *Lecture Notes in Computer Science*, pages 280–292. Springer, 2001.

- [8] Zhen Liu, Srinivasan Parthasarathy, Anand Ranganathan, and Hao Yang. A generic flow algorithm for shared filter ordering problems. In Maurizio Lenzerini and Domenico Lembo, editors, *PODS*, pages 79–88. ACM, 2008.
- [9] S. Salloum. *Optimal testing algorithms for symmetric coherent systems*. PhD thesis, Univ. of Southern California, 1979.
- [10] S. Salloum and M. A. Breuer. An optimum testing algorithm for some symmetric coherent systems. *J. Mathematical Analysis and Applications*, 101:170–194, 1984.
- [11] S. Salloum and M.A. Breuer. Fast optimal diagnosis procedures for k-out-of-n:g systems. *Reliability, IEEE Transactions on*, 46(2):283 –290, jun 1997.
- [12] Tonguç Ünlüyurt. Sequential testing of complex systems: a review. *Discrete Applied Mathematics*, 142(1-3):189 – 205, 2004. Boolean and Pseudo-Boolean Functions.